



ComPASS  
SciDAC-3

---

# ComPASS

## Present and Future Computing Requirements

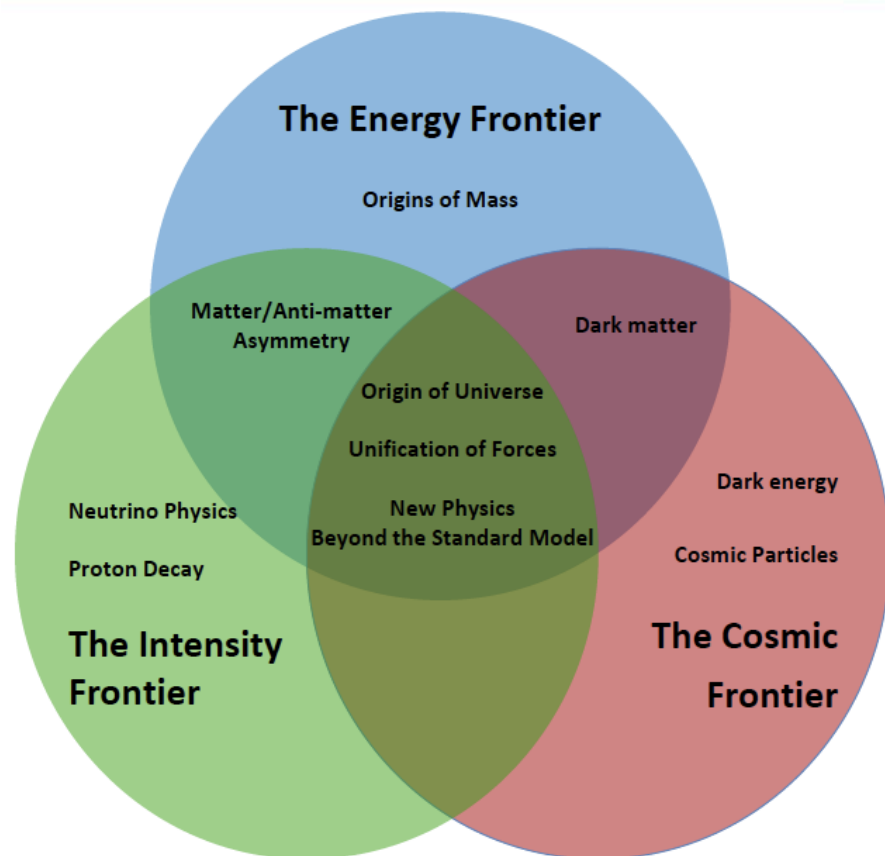
Panagiotis Spentzouris (Fermilab)  
for the ComPASS collaboration

*NERSC BER Requirements for 2017  
September 11-12, 2012  
Rockville, MD*



# Accelerators for High Energy Physics

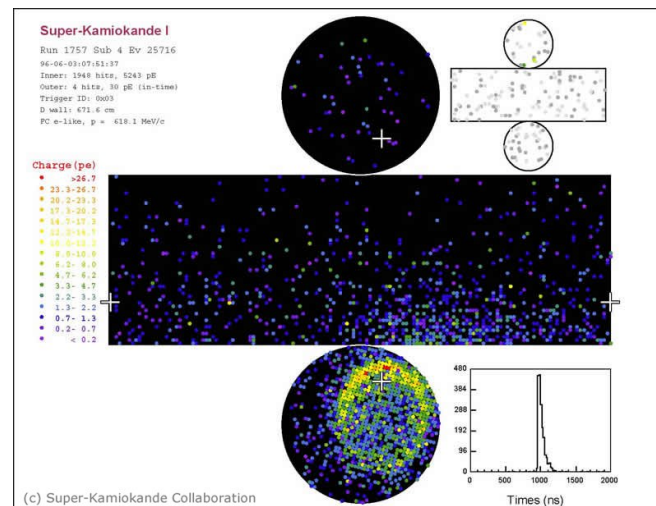
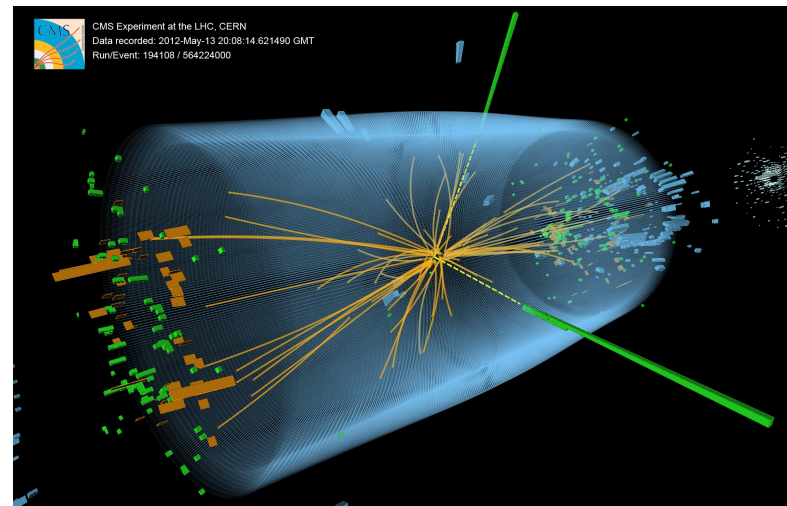
- At the Energy Frontier, high-energy particle beam collisions seek to uncover new phenomena
  - the origin of mass, the nature of dark matter, extra dimensions of space.
- At the Intensity Frontier, high-flux beams enable exploration of
  - neutrino interactions, to answer questions about the origins of the universe, matter-antimatter asymmetry, force unification.
  - rare processes, to open a doorway to realms to ultra-high energies, close to the unification scale
- Particle accelerators indirectly support the cosmic frontier by providing measurements of relevant physics processes





# Where we are today

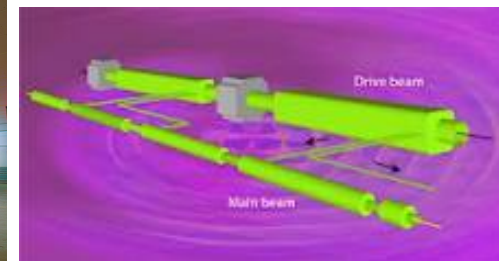
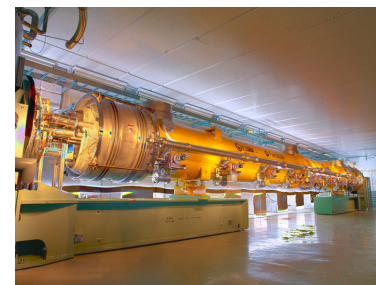
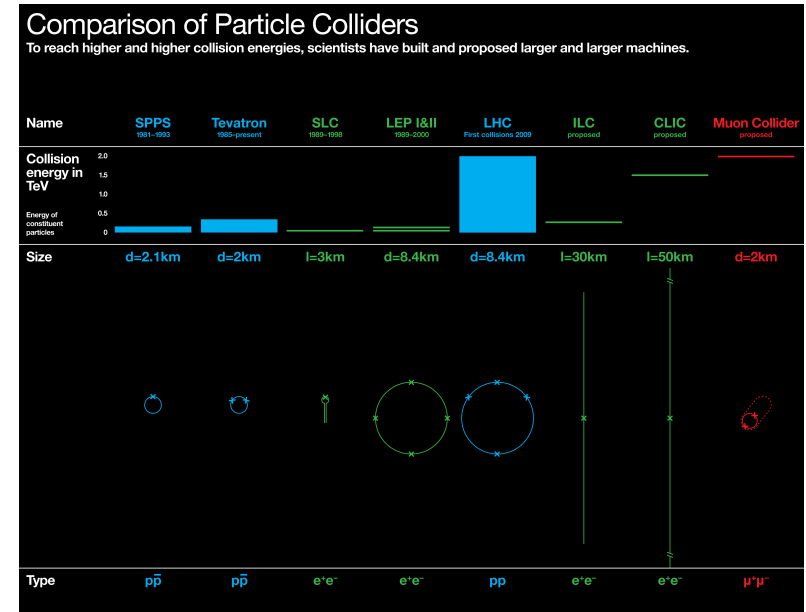
- Discovery of the Higgs particle, responsible for electroweak symmetry breaking and the mass of elementary particles
  - No physics beyond the “Standard Model” (SM) of HEP has been observed
- Neutrinos oscillate, thus have mass
  - No answers on mass hierarchy or symmetry properties





# Where we would like to be (Energy Frontier)

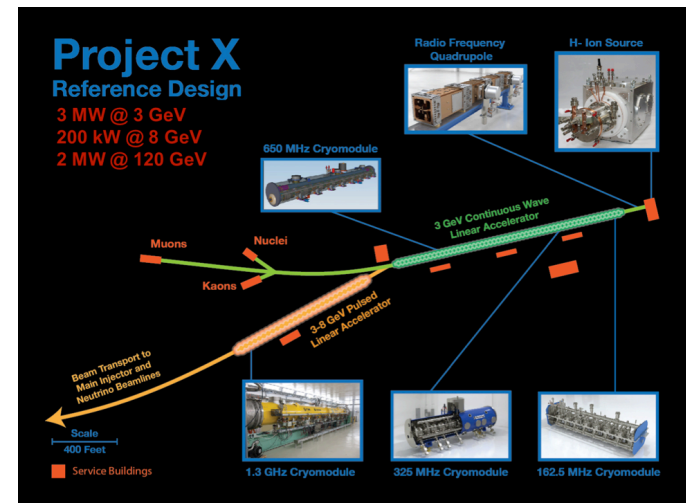
- A dedicated accelerator will be necessary to study Higgs properties
  - Is it a “Standard Model” Higgs?
- “Higgs Factory” candidate: lepton collider
- A great challenge for accelerator science!
  - Develop techniques, technologies and materials to achieve higher acceleration gradients
    - dielectric and plasma wave structures, beam cooling
  - Optimize existing technologies
    - Superconducting rf cavities
  - Optimize and test designs
    - CLIC, Muon Collider





# Where we would like to be (Intensity Frontier)

- A high-intensity proton accelerator to drive
  - long-baseline neutrino oscillation experiments
    - Mass hierarchy, matter-antimatter asymmetry, oscillation parameters
  - muon, kaon experiments
    - Physics beyond the SM
- Staged approach at Fermilab
  - Improvements of existing machines
  - New linear accelerator: Project-X
- A great challenge for accelerator science!
  - Controlling instabilities to minimize beam losses is essential
    - Self-fields, wakefields, interaction with materials, geometry and long term tracking accuracy

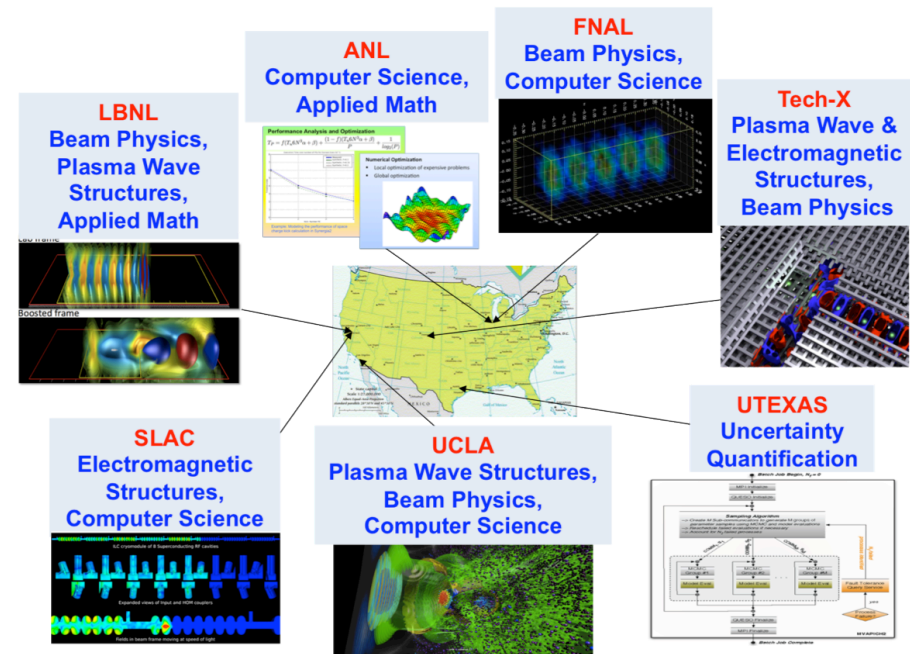




# Advanced Computation for HEP Accelerator Science and Technology

- To enable scientific discovery in HEP, high-fidelity simulations are necessary to develop new designs, concepts and technologies for particle accelerators
- Under SciDAC3, ComPASS develops and deploys state-of-the-art accelerator modeling tools that utilize
  - the most advanced algorithms on the latest most powerful supercomputers
  - cutting-edge non-linear parameter optimization and uncertainty quantification methods.

## The ComPASS collaboration



Community Project for Accelerator  
Science and Simulation (ComPASS)



# ComPASS Methods and Tools

- A comprehensive set of codes that incorporate state-of-the-art field solvers
  - Electrostatic: multigrid (*Synergia*, *Warp-FastMATH*); AMR multigrid (*Warp-FastMATH*)
  - Electrostatic: spectral (*Synergia*)
  - Electromagnetic: finite element direct and hybrid (*ACE3P-FastMATH*)
  - Electromagnetic: extended stencil finite-difference (*Osiris*, *Vorpal*, *Warp-FastTMATH*); AMR finite-difference (*Warp-FastMATH*)
  - Quasi-static: spectral (*QuickPIC*)
- Software Applications and Tools
  - Chombo, FFTW, HDF5, LAPACK, METIS, MPI, MPI/IO, MUMPS, ScaLAPACK, SuperLU, TAU, TAO, Trillinos
  - Shared library support is very important
  - C++, Fortran90, Python
  - For analysis we use ParaView, Python, R Language, ROOT, VisIt.

**The ComPASS toolkit:**  
ACE3P, Osiris, QuickPIC,  
Synergia, Vorpal, Warp



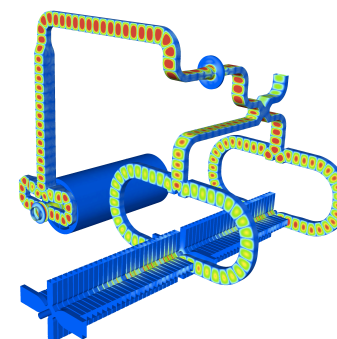
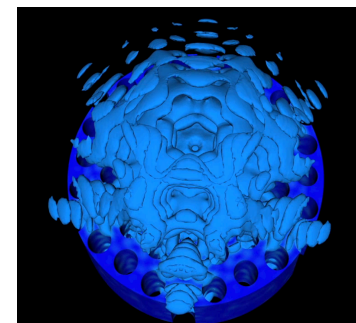
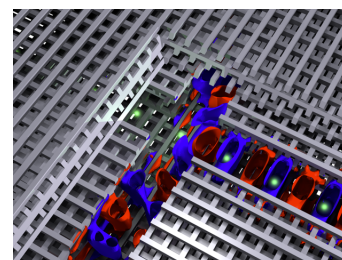
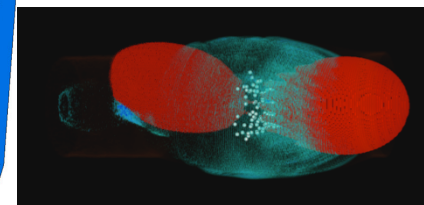
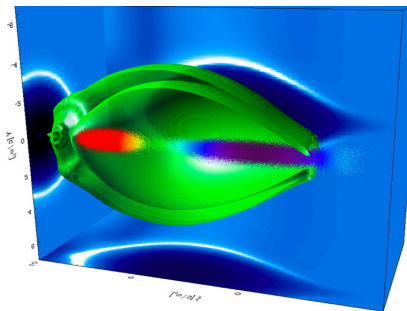
# ComPASS SciDAC3 applications

- Support the development and study of new technologies for smaller and possibly cheaper energy frontier accelerators:
  - accelerators based on standard technology are limited by the metallic electrical breakdown limit of 50-100 MV/m
  - dielectric laser accelerators: a laser propagating through a dielectric lattice can generate electric fields of few GV/m
  - plasma based acceleration: a driver beam (laser/particles) propagating through a plasma creates a wake with accelerating gradients exceeding 50 GV/m.
- Focus on plasma and dielectric R&D and optimization of conventional technology applications
- Support the design and optimization of high-intensity proton accelerators to minimize beam losses that cause radiation damage. Modeling of
  - many (all) beam bunches in circular machines and their coupling through impedance and wakefields
  - beam self-charge and instabilities caused by beam-matter interactions
  - field non-linearities and accelerator geometry (apertures, positions, fields)
- Focus on Fermilab existing proton source improvements and Project-X



# Energy Frontier Objectives

- Plasma-based acceleration:
  - support the BELLA (laser) and FACET (beam) experimental programs
  - develop techniques to improve beam quality
  - study controlled electron beam injection
  - improve staging for future lepton collider concepts.
- Dielectric laser acceleration:
  - design efficient power couplers between optical fiber and accelerator structure
  - explore wakefield effects and associated break-ups for different topologies
  - design structures able to accelerate high quality beams
- High Gradient acceleration:
  - understand wakefields in the Power Extraction and Transfer Structure (PETS) system of CLIC

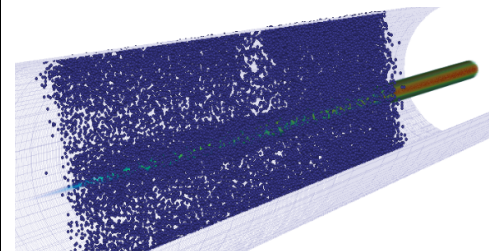
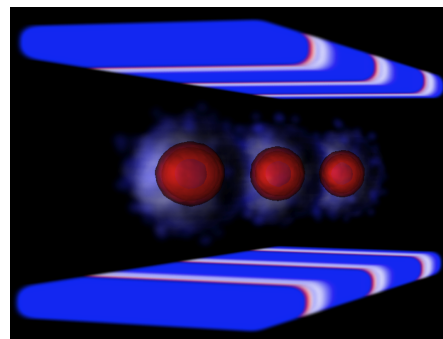




# Intensity Frontier Objectives

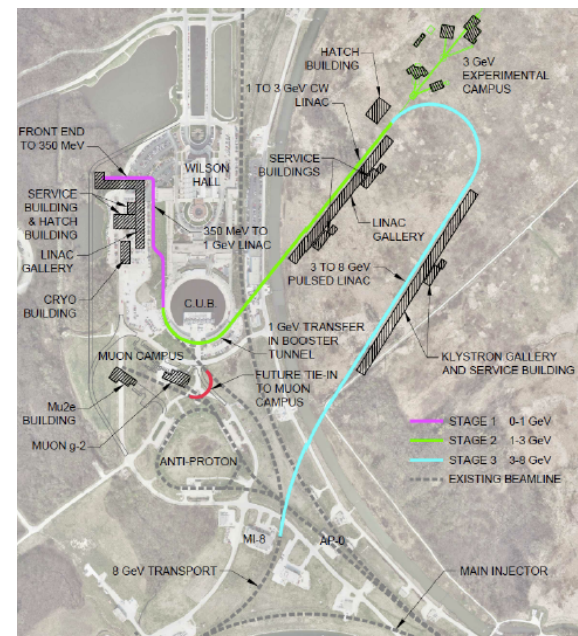
## ■ Fermilab proton source upgrades for the Neutrino and Muon Programs

- Booster synchrotron: instability control for beam quality and loss minimization (targeting 50% increase of beam flux)
- Main Injector (MI) synchrotron: instability mitigation and loss minimization (targeting 100% increase of beam flux)



## ■ Project-X: support staging

- Study wakefields for the first stage of the linac
- Model experiments of electron-cloud effects in the MI, currently underway
- Study mitigation techniques to control losses in MI due to self-fields, wakefields, and electron-clouds



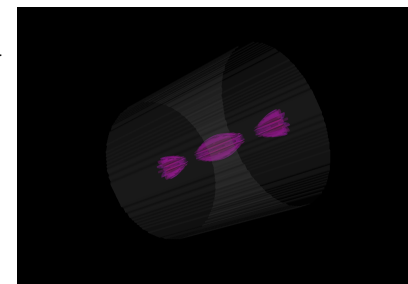
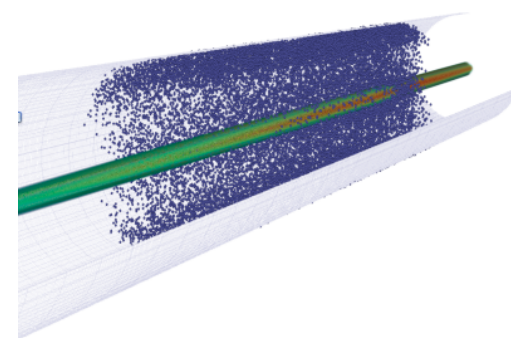


# ComPASS HPC resources

- The ComPASS NERSC repository is used for SciDAC supported code development and applications (6M hours in 2012). ComPASS codes are also used for applications in the NERSC repositories discussed by Geddes, Ko, and Tsung.
- ComPASS researchers utilize ALCF resources (5M hours, becoming 80M hours in 2013), and OLCF.
- Here we discuss HPC resource requirements for accelerator science related to Intensity Frontier deliverables.



# Modeling a High-Intensity Accelerator



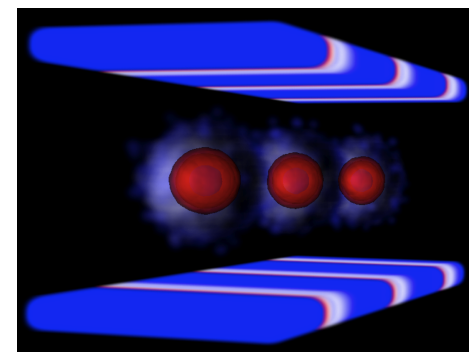
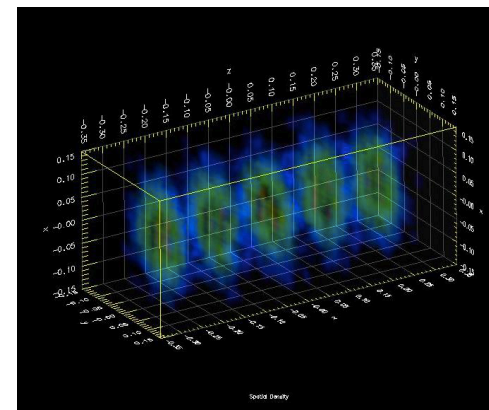
## ■ Wide range of scales:

- accelerator complex ( $10^3\text{m}$ )  $\rightarrow$  EM wavelength ( $10^2\text{-}10\text{ m}$ )  $\rightarrow$  component ( $10\text{-}1\text{ m}$ )  $\rightarrow$  particle bunch ( $10^{-3}\text{ m}$ )
- Need to correctly model intensity dependent effects to identify and mitigate potential problems due to instabilities that increase beam loss and the accelerator lattice elements (fields, apertures)



# Modeling a High-Intensity Accelerator

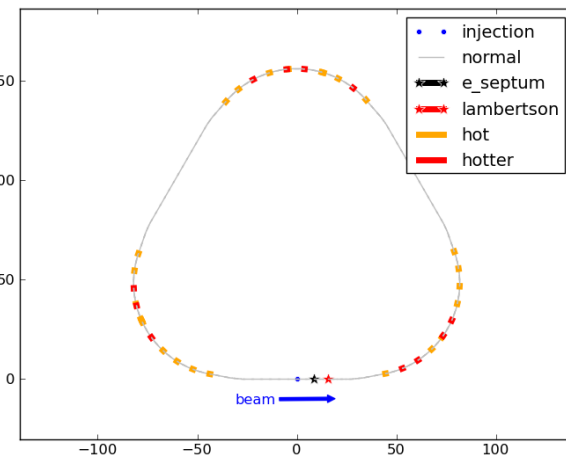
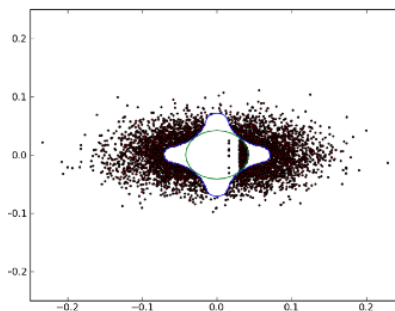
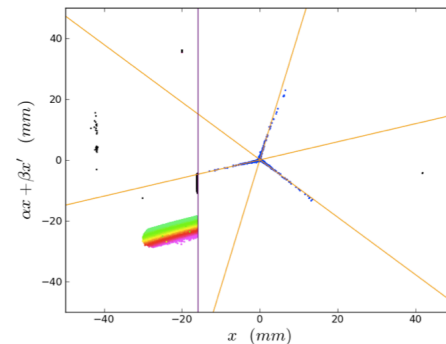
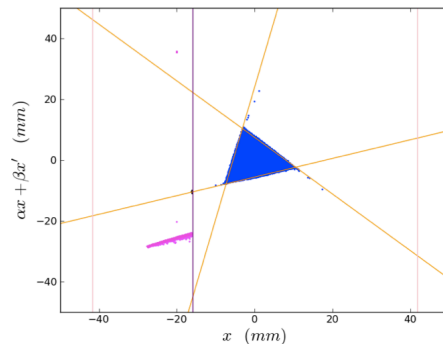
- Beam particles form tightly packed bunches; usually many bunches circulate the machine
- Individual particle motion is controlled by accelerator elements (magnets, rf cavities)
- Within a beam bunch intensity-dependent effects affect particle motion
  - space charge (repulsion between beam particles)
  - wakefields induced in the accelerator structures by the beam
  - electron clouds, generated on the structure walls and amplified by the passing beam
- Wakefield effects occur both within a single bunch (head-tail instabilities), and/or between bunches (coherent bunch motion)
- Intensity-dependent effects lead to beam loss by exciting resonances formed by nonlinearities due to imperfections in the machine structures





# Intensity Frontier: Mu2e extraction design

- Synergia full extraction simulation for the Muon to Electron (Mu2e) Fermilab proposed experiment
  - Single bunch, 1M macroparticles
  - 26k turns, 240 3D solves/turn
  - Realistic apertures and non-linear fields
  - Quantitative loss predictions
  - Parameter optimization runs (septum location, magnet ramp)
- Results contributed to CD1 approval for the experiment





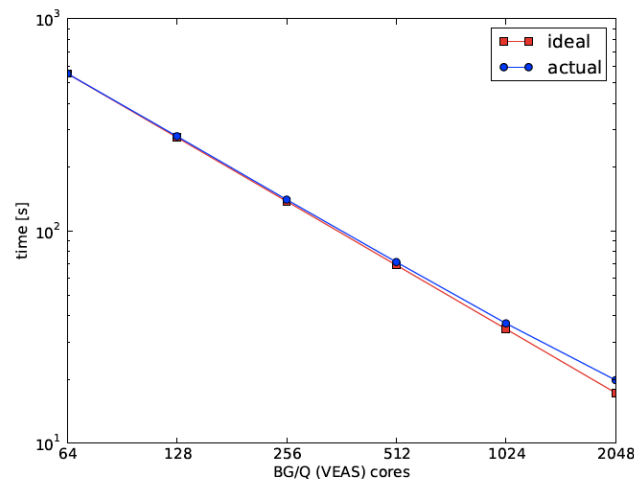
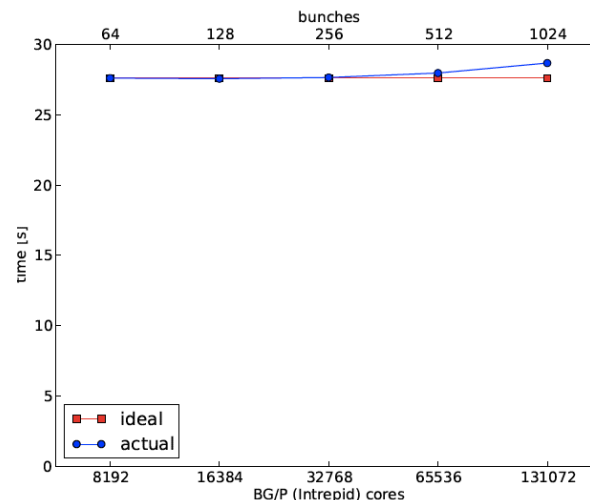
# Current HPC usage

- Hours: 6 M/year (NERSC), see also Geddes, Ko, and Tsung
  - ALCF (5M 2012 going to 80M 2013) and OLCF
- Cores:
  - 3D PIC average ~ 16 kcore
  - Mu2e example ~ 2kcore (Intensity Frontier SciDAC2 example, relevant to SciDAC3 and 2017 discussion)
- 10's of simultaneous runs
- Scaling I/O & queue limited
- Archival data: ~TB
- Run data ~TB/run



# Goal: Main Injector for Project-X

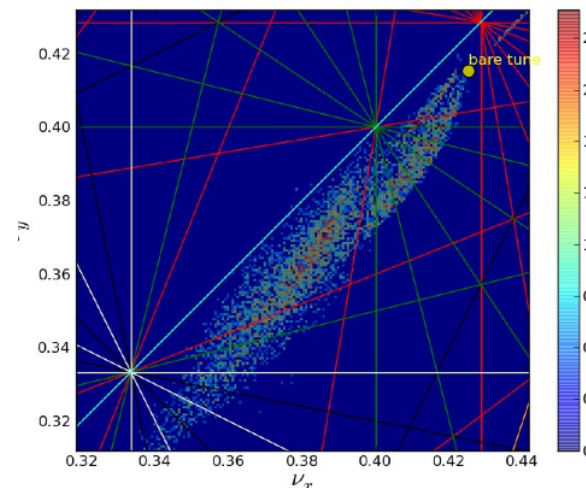
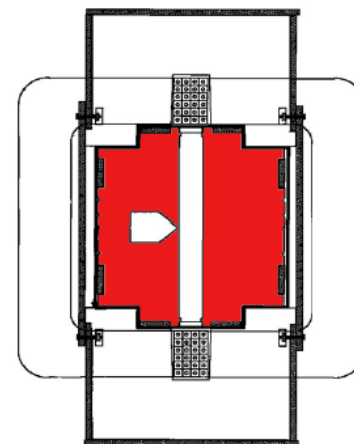
- Produce a map of expected losses as a function of machine configuration in tune space, find optimal operating point
  - Single bunch, space-charge and impedance
  - Include uncertainties in measured multipole field errors, realistic apertures
- Loss prediction for all 588 bunches (coupled by impedance), re-optimize operating point (scan)
- The parameter optimization requires loosely coupled groups of ~100 jobs consisting of 1024-2048 core tightly-coupled calculations
- The multi-bunch runs 588 bunches at 128-512 cores/bunch





# I/O needs

- In order to make accurate predictions of beam loss it is crucial to have a realistic model of the accelerator components
  - Positions, apertures, field errors
- Detailed particle tracking is necessary to analyze the effects of physical apertures.
  - With “playback” capabilities, in order to identify problematic areas of phase-space
  - This is a paradigm shift for such applications which makes efficient parallel I/O very important for future systems
    - Without sub-sampling or introducing selection algorithms x1000 of current output (write all particles for every step)





# HPC requirements 2017

- Hours: 85 Million/year
  - Assuming continuation of current model (additional INCITE allocations at ALCF, NERSC, OLCF for specific applications)
  - Driven by accelerator parameter optimization (“design optimization” runs) & multi-physics, multi-scale models
- Cores: weak scaling dominant
  - 75-300k, problem and network dependent
  - Up to 100 of simultaneous runs at few kCore for parameter scans
- Memory ~1 GB/core



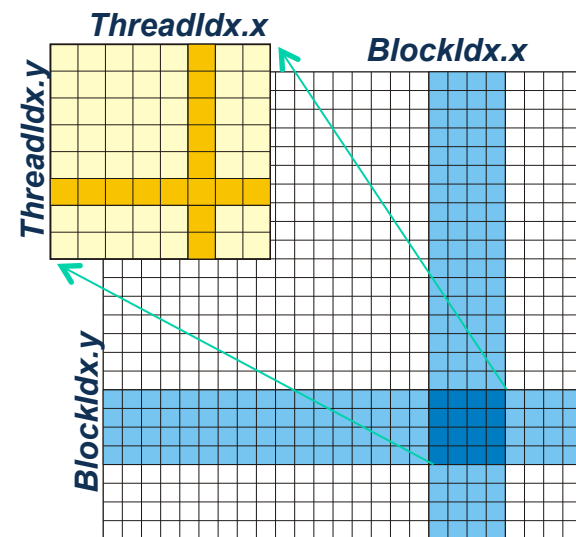
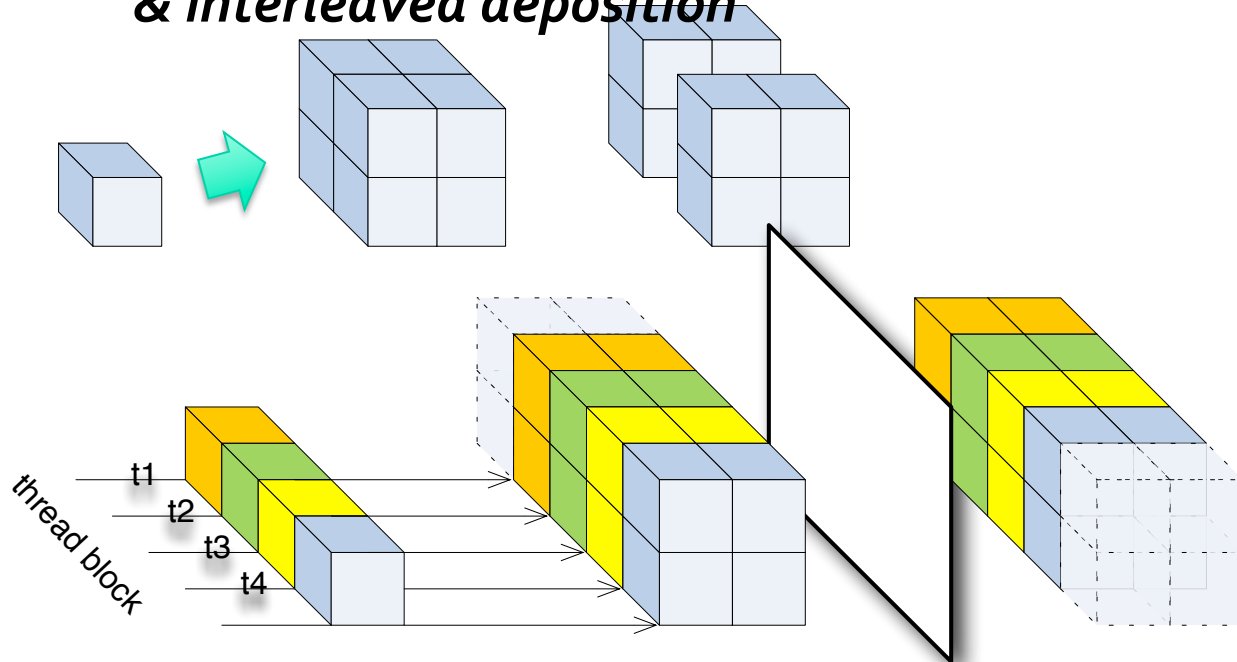
# New Architectures

- We are already working on algorithmic and code development that will enable our codes to perform on new architectures (“lightweight” processors, accelerators, or hybrid configurations).
- Our strategy is to abstract and parameterize data structures so that are portable and enable efficient flow of data to a large number of processing units in order to maintain performance.
  - Have already ported subset of solvers and PIC infrastructure on the GPU (Synergia, UPIC, VORPAL)
  - New algorithms based on hybrid MPI+OpenMP (Synergia)
    - Both efforts funded by ASCR @ FNAL
- We will need sufficient notice and specs of the chosen new NERSC architecture and a test system well in advance of deploying the new production machine to fully take advantage of this work.



# Synergia on the GPU

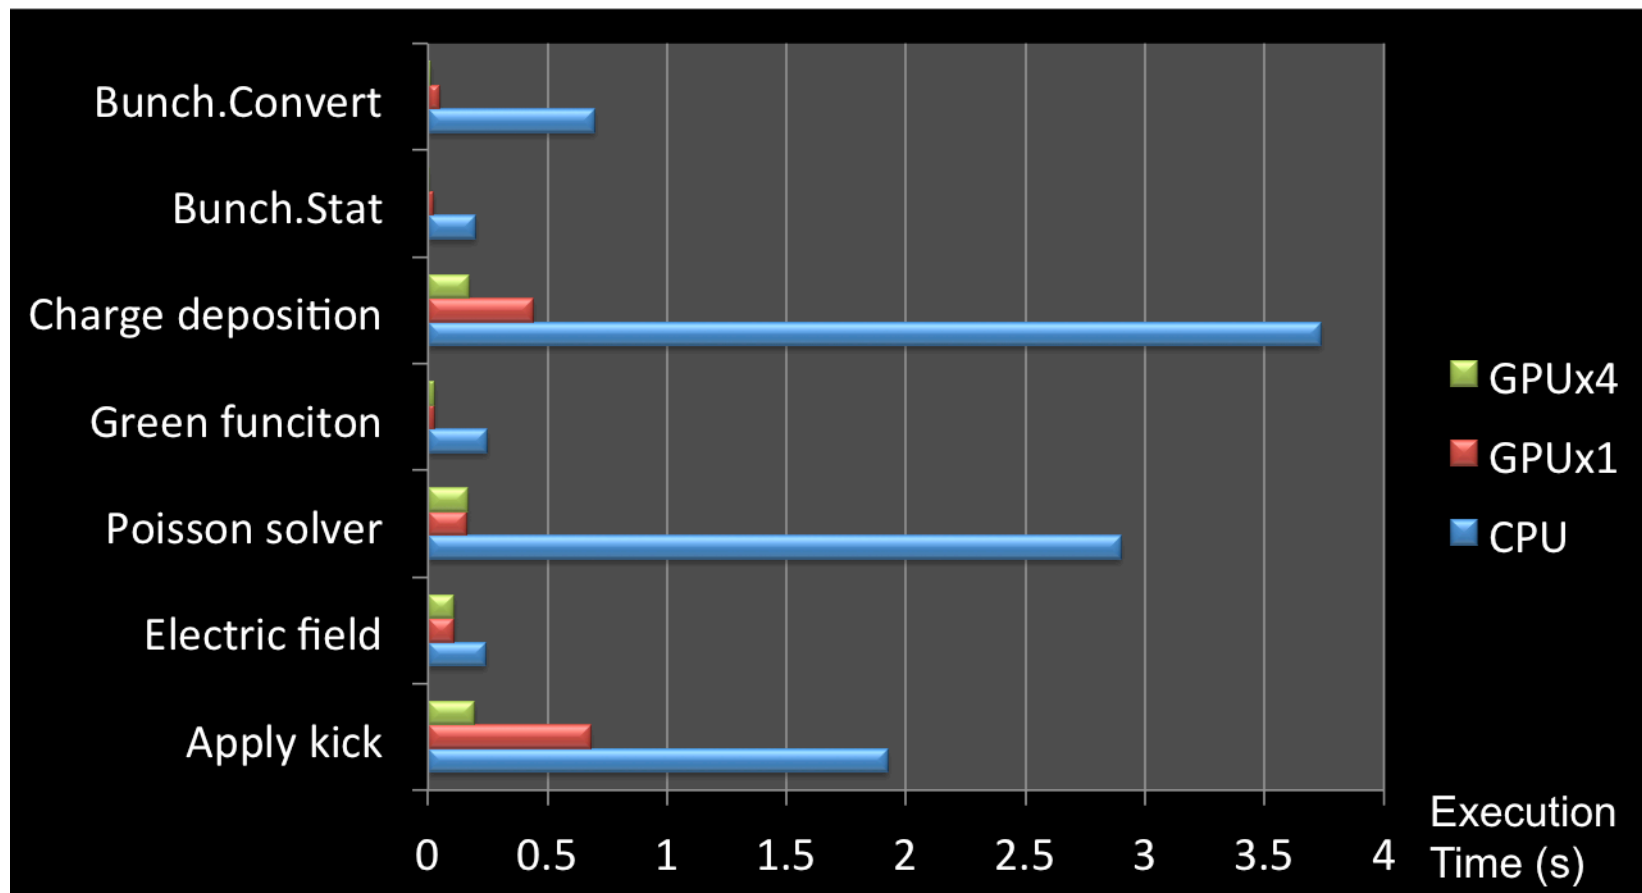
*Sort particles to cells  
& interleaved deposition*



- Use shared memory for field deposition
- One thread per cell for particle kicks



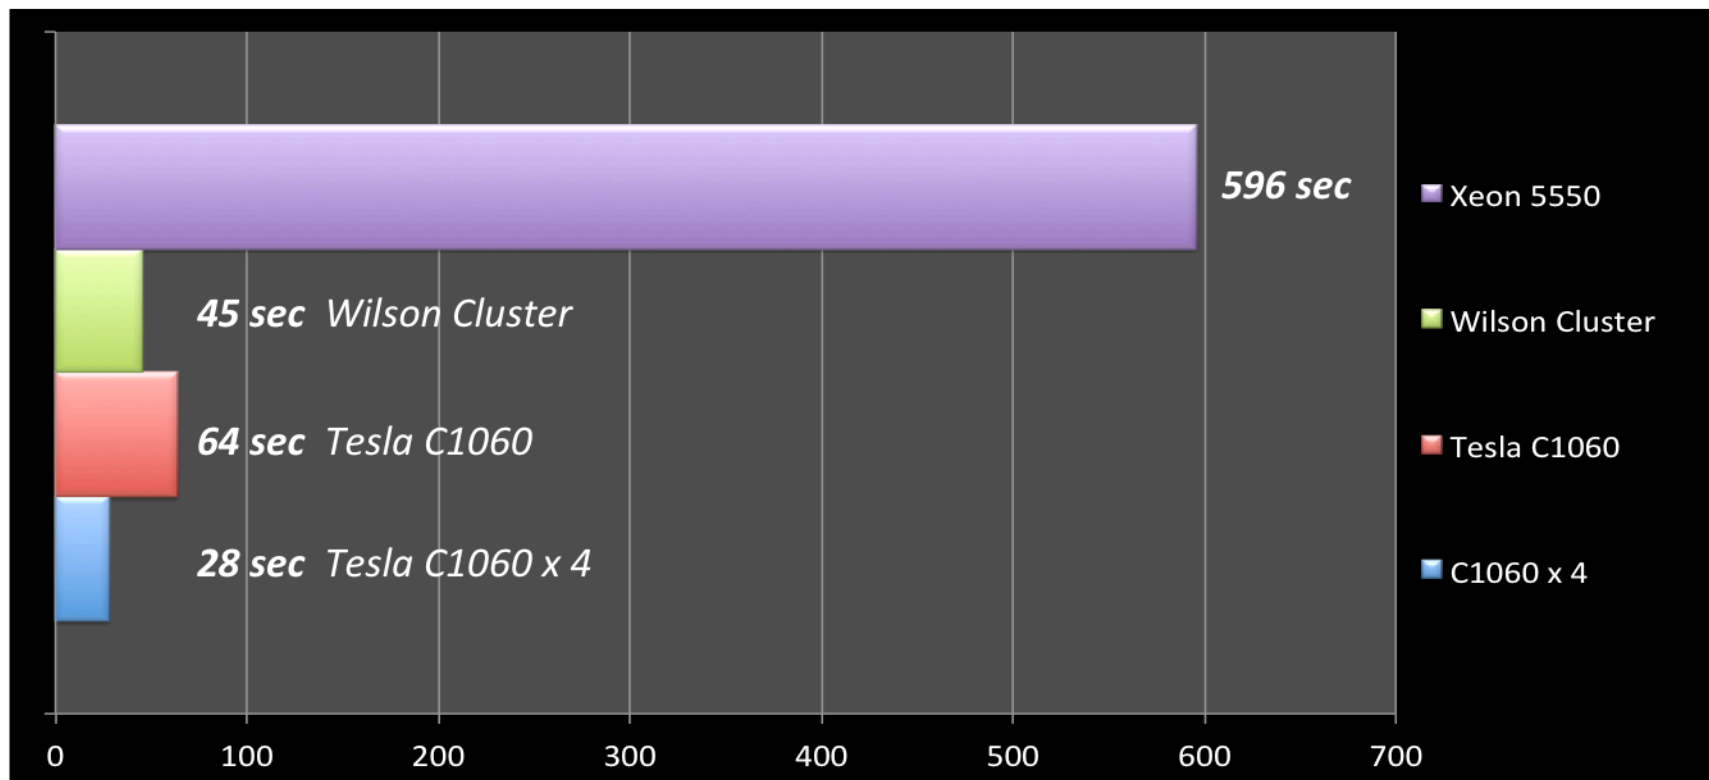
# Synergia on the GPU



1. Intel Xeon X5550, single process @ 2.67GHz;
2. NVidia Tesla C1060, 30 streaming multi-processors @ 1.30GHz in a single GPU
3. Nvidia Tesla C1060 x 4



# Example: Synergia on the GPU



1. Intel Xeon X5550, single process @ 2.67GHz;
2. Fermilab Wilson Xeon Cluster, dual Xeon X5650 2.67GHz nodes. 16 nodes / 128 cores used
3. NVidia Tesla C1060, 30 streaming multi-processors @ 1.30GHz in a single GPU
4. NVidia Tesla C1060 x 4



# Summary

	Used at NERSC in 2012	Needed at NERSC in 2017
Computational Hours	6M	85M
Typical number of cores* used for production runs	16K	75K
Maximum number of cores* that can be used for production runs	64K	300K
Data read and written per run	1TB	10TB
Maximum I/O bandwidth	0.3GB/sec	3GB/sec
Percent of runtime for I/O	20	20
Shared filesystem space	10TB	100TB
Archival data	1TB	500TB
Memory per <del>node</del> core	1GB	GB
Aggregate memory	1TB	10TB